EDGE INTELLIGENCE – UNLEASHING THE FULL POTENTIAL OF VIRTUAL CONTACT CENTERS

A WHITE PAPER ON THE BENEFITS OF EDGE – ARCHITECTED CONTACT CENTER APPLICATIONS

BY: SPANLINK COMMUNICATIONS

SPANLINK™

The success of IP and the Internet has taught the IT world some important lessons. SIP, the Session Initiation Protocol, has become the de facto signaling standard for IP communications, including voice, video, instant messaging and other real-time media. SIP, which is supported by equipment manufacturers, software houses and service providers worldwide, favors intelligent communications devices at the network edge – while the bulk of the complexity and large servers can be hosted anywhere in redundant data centers – near the experts that can manage them.

This "edge-oriented" model yields more scalable, more robust, and more resilient business services. Now those lessons are transforming enterprise call centers. Economical voice over IP is replacing traditional TDM-based voice. Distributed "virtual" call centers are making single-site facilities obsolete. And IP-enabled voice-data convergence is enabling powerful new applications.

But the transformation is still incomplete. Although big national call centers are giving way to the virtualization of human resources at regional and home locations, the call center **applications** still remain highly centralized. All business logic still executes on central servers, and end-user PCs still act as little more than dumb terminals.

As enterprises update their voice systems, replacing TDM "silos" with a distributed VoIP infrastructure, they would be wise to review their application architectures as well. By moving certain key applications to the network edge, enterprises can eliminate the server bloat, excessive bandwidth costs, scalability barriers, wasted resources and productivity constraints imposed by today's centralized systems. Indeed, the full promise of IP communications demands a new application architecture: an edge-oriented model that puts business logic where it is most efficient and effective, uses the full resources of the distributed infrastructure, and—like grid computing—leverages the aggregate processing power and storage capacity of end-user PCs.

This white paper introduces Edge Intelligence, an innovative approach to call center applications. By leveraging both the flexibility of IP communications and the power of distributed computing, Edge Intelligence brings several valuable benefits to enterprise call centers:

- Faster, easier application development,
- Increased end-user productivity,
- More reliable system performance,
- Expanded system capacity, flexibility and scalability,
- Reduced capital and operating expense.

HALFWAY THERE: THE MOVE TO VIRTUAL CALL CENTERS

The advantages of IP communications have not gone unnoticed by enterprises. Call centers of all types—travel reservations, customer support, catalog sales, etc.—are using VoIP and IP telephony to reduce costs through virtualization. Instead of bringing everyone together—reservation agents, subject matter experts, sales specialists, supervisors, etc.—under one large roof, enterprises are setting up smaller regional call centers. Many companies are even allowing call center employees to work from home and other virtual company locations.

In fact, anyone with a personal computer and broadband Internet access can be part of a virtual call center. IP networking of both voice and data ties distant locations together without creating separate site-by-site silos. Dynamic routing makes this "virtualness" transparent to callers and makes distributed workers as efficient and productive as if they were all in one room.

The economic benefits of IP-based virtual call centers are clear: reduced bricks-and-mortar costs, smaller phone bills, and more efficient response to shifting demand. Rather than maintaining around-the-clock shifts at a central site, the call center can move through time zones with the sun. And there is no need to maintain facilities for peak traffic loads, since an enterprise can easily add agents and increase call center capacity on demand. Instead of waiting for the phone company to provision leased lines or ISDN service, temporary workers simply join the virtual call center over the Internet.

However, while modern call centers may be reaping the benefits brought by virtualizing people and devices, their application architectures still reflect the old centralized model.

By continuing to host all application logic at a central site, enterprises are failing to realize the full economic and functional benefits of IP communications.

THE CENTRALIZED MODEL: THE CURRENT STATE OF PRACTICE

The current state of practice for IP-based call centers depends on centralized application architecture reminiscent (Figure 1). The lowest tier of the architecture is a central "farm" of backend servers holding customer databases, flight schedules, product catalogs, inventory records or whatever data the applications require. The next tier, usually at the same location, is an array of servers running the applications that execute business

processes: make or change a reservation, place an order, retrieve billing information, search a parts list, record a conversation, and so on. The topmost tier comprises end-user PCs and workstations in major call centers, regional locations and workers' homes. "Thin" client software on each end-user PC accepts keystrokes and mouseclicks and paints the computer screen (Figure 2). Any information that a user enters is sent to the central servers for processing. Any output—including input forms and command menus—is generated by the central servers and sent to the PC for display.

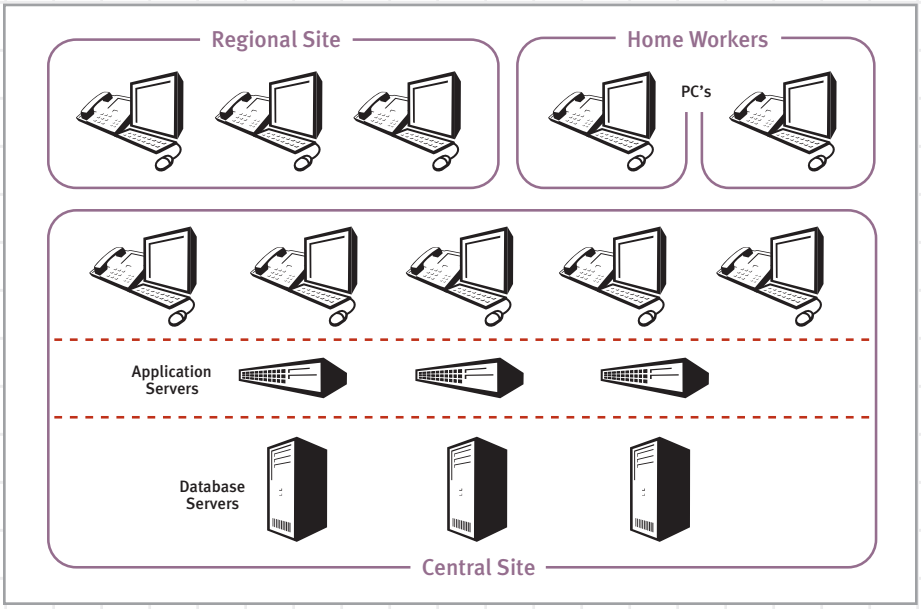

Figure 1. Virtual call centers employ a three-tiered architecture.

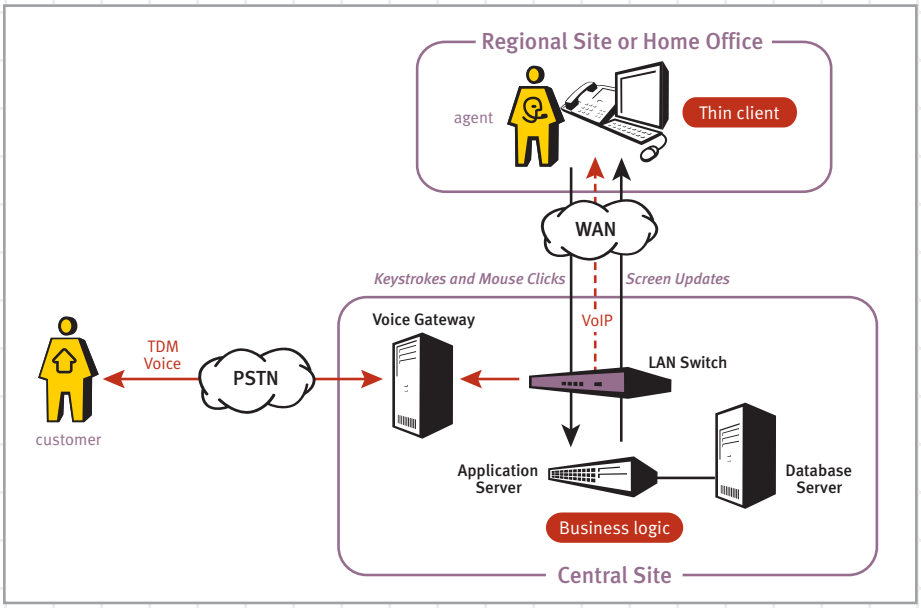

Figure 2. With centralized architecture, all business logic resides on central application servers.

Centralized architecture has a few advantages. Databases are easy to maintain since they are co-located. Software distribution is relatively simple since end-user PCs rely only on thin clients and changes to business logic can be applied centrally. But centralized architecture also has major shortcomings:

> **Server bloat** – Large numbers of application servers are needed at the central site, since all business logic is executed there. The complex multi-threading required to support hundreds or thousands of simultaneous transactions causes server software to expand even further. And since the loss of an application server can affect hundreds or thousands of users, the enterprise must deploy redundant servers. The cost to deploy and maintain so many application servers can be staggering.

> **High bandwidth costs** – With only thin clients at the network edge, every transaction requires several exchanges of data between the end-user PC and the central site. Selected voice traffic also travels to the central site for recording. The only way to ensure responsive performance, therefore, is to provision—and pay for—large amounts of high-priority network bandwidth.

> **Diminished productivity** – To offset the complexity of multi-threading and to limit bandwidth consumption, developers must "dumb down" applications and user interfaces, resulting in diminished end-user productivity and a steeper learning curve for new employees.

> **Reliability challenges** – Multi-threaded applications are inherently more complex and error prone than single-threaded software. Moreover, the loss of the central site can mean the loss of all service. The only way to ensure continuous service is to back up the central site with a duplicate server farm at another

location—an especially expensive proposition due to server bloat.

> **Scaling challenges** – For even a temporary uptick in demand, application capacity can be increased only by adding costly servers, plus additional backup and network capacity.

> **Wasted resources** – While central server farms struggle to keep up with application loading, most of the computing and storage capacity of end-user PCs sits idle.

Centralized application architecture, therefore, is not the best approach for every business. Although it does allow the distribution of human resources to create virtual call centers, it fails to exploit the full capabilities of IP communications. A benefit of IP is its ability to access resources anywhere in the infrastructure, not just at a central location. So with IP, business requirements—instead of resource constraints or connectivity limitations—can drive application architecture.

EDGE INTELLIGENCE: UNLEASHING THE FULL POTENTIAL

Edge Intelligence, a new model for IP-based call center applications, builds on lessons learned from the Internet. By moving business logic closer to the information and events that drive it and by mobilizing the full computing power and storage capacity of end-user PCs, Edge Intelligence eliminates many of the shortcomings of the centralized approach.

Like the centralized model, Edge Intelligence comprises three tiers or layers:

> Centrally-located database servers at the bottom,

> Common business logic servers in the middle,

> End-user PCs and other intelligent devices at the top.

But application intelligence is distributed differently (Figure 3). Instead of moving information to the business logic, business logic moves to the information. Where it makes sense, powerful applications replace thin clients in end-user PCs. For example, instead of shipping real-time voice to central servers for recording and analysis, PC-based applications record and analyze voice calls at the edge.

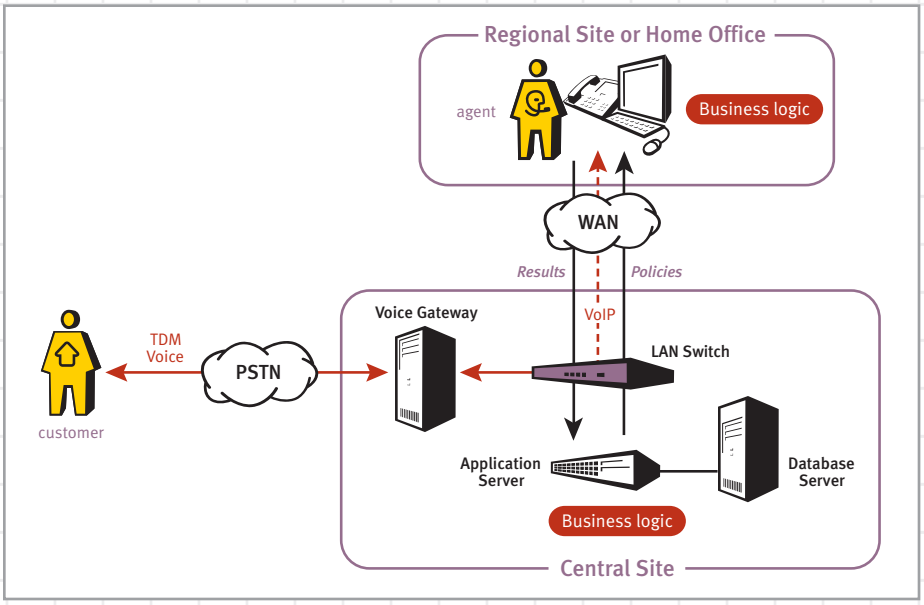

Figure 3. With centralized architecture, all business logic resides on central application servers.

Like centralized architecture, Edge Intelligence still handles hundreds or thousands of concurrent processes. But much of the workflow is handled as independent streams on distributed PCs, instead of multi-threaded processes on central servers. The call center still relies on central servers for administrative functions like user maintenance and authoring of workflow rules. Central databases still hold information like evaluation scores associated with recorded contacts and call data like date, time, ANI, etc. But the central server layers are much thinner, since there's less centralized work to do.

Of course, one size does not fit all. The amount of business logic that can move to the edge depends on the nature of the applications, where information enters the system, and where events occur. But for call centers that rely on both voice and data interaction at user workstations, Edge Intelligence promises major benefits:

> **Reduced server bloat** – Software on central servers is more compact because much of the business logic is handled by edge devices and there is less multi-threading. The enterprise enjoys a reduction in both capital and operating expense since it requires fewer central servers.

> **Network cost savings** – Since application processing and even voice recording occur at the edge, without the back-and-forth of keystrokes and screen updates, network loading and network costs decrease dramatically.

> **Increased user productivity** – The learning curve flattens and user productivity improves because single-threaded applications on PCs can support more complex interactions

and friendlier user interfaces than the multi-threaded centralized model.

> **Greater reliability** – It is easier to create reliable applications for single-user PCs than for multi-user central servers. And since edge devices operate independently of each other, an application "crash" affects only one user. Relieved of much multi-threaded complexity, central servers are also less error prone. And in many cases, applications on remote PCs can continue to function—and end-users can complete transactions—during a network or central-site outage.

> **More responsive, more affordable scalability** – Edge Intelligence helps call centers respond quickly and economically to fluctuating demand. With application processing moved to the edge, infrastructure capacity can grow by adding low-cost PCs instead of expensive servers. And there is no need to rent extra office space during peak seasons since many employees can connect from home.

> **More efficient resource utilization** – With Edge Intelligence, the processing capacity and capital expense of end-user PCs is no longer wasted. Instead of emulating dumb terminals and leaving most of their CPU power and disk space idle, PCs become the workhorses of the call center. The aggregate processing power and storage capacity of the virtual call center increases dramatically while the total cost of computing resources may actually shrink.

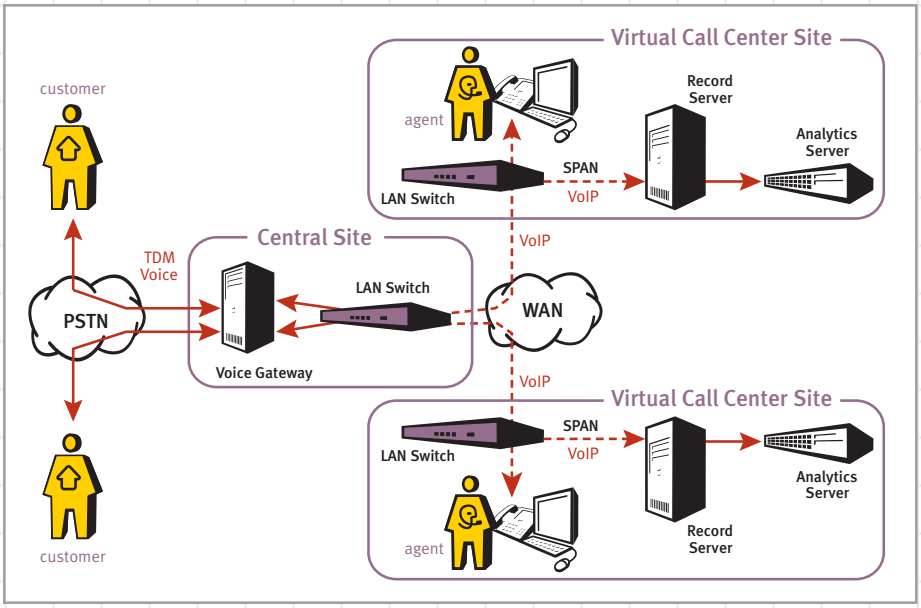

Figure 4. Traditional solutions copy VoIP conversations to site-based record servers via SPAN ports on local LAN switches.

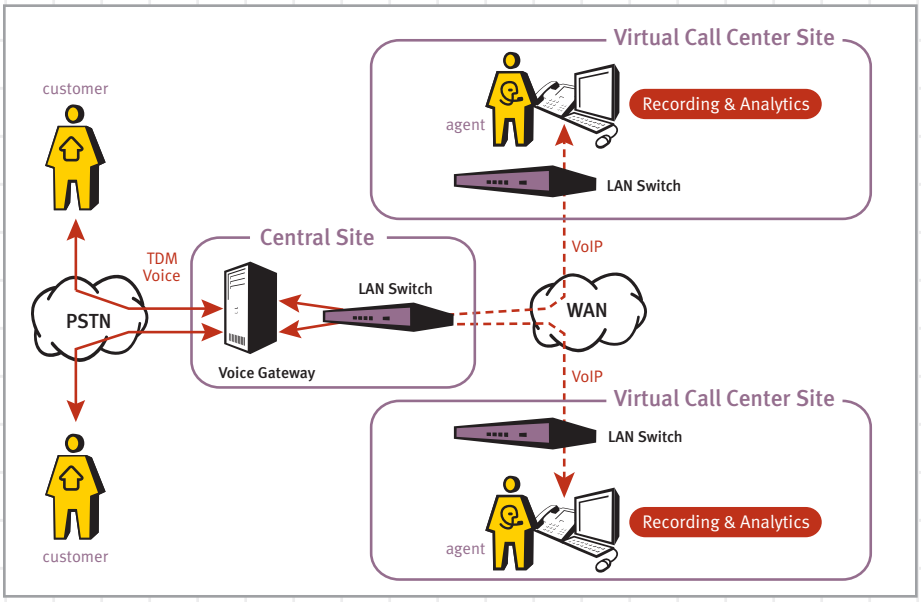

Figure 5. The Edge Intelligence solution records and analyzes live conversations at each PC.

Of course, compared to the centralized model, Edge Intelligence adds some operational challenges. Call center operations staff have to maintain diverse applications on large numbers of remote PCs and manage policies that apply to large numbers of remote users. But these operational challenges can be minimized by configuring remote PCs to check with the central site on start-up and automatically download any software or policy updates.

IMMEDIATE BENEFITS: TWO EDGE INTELLIGENCE EXAMPLES

Perhaps the best way to understand Edge Intelligence is through examples. Voice recording and analytics and workflow automation are two areas where Edge Intelligence delivers immediate benefits.

Voice Recording and Analytics

For quality control, regulatory compliance and other reasons, call centers often record and analyze employees' conversations with callers. If those conversations are recorded and

analyzed by dedicated servers at central or regional sites, the resulting bandwidth costs can be prohibitive. By moving those processes to the network edge, the call center can realize huge reductions in bandwidth utilization and associated costs.

The traditional approach to voice recording and analytics records selected conversations on site-based record servers (Figure 4). A central voice gateway takes customer calls from the public switched telephone network (PSTN), converts them into VoIP packet streams, and forwards them to virtual call center workers over the enterprise wide-area network (WAN). At each site, a LAN switch copies the VoIP traffic to a local record server that reassembles and saves the packet streams as audio files. The audio files may then travel to special analytics servers for content processing.

This method requires expensive record servers—perhaps even redundant record servers—at every site. Configuration of these servers can be challenging and can complicate network design. Moreover, to capture live VoIP streams, the record servers must attach to Switch Port ANalyzer (SPAN) ports on the LAN switches, reducing SPAN port availability for network monitoring and other important applications. (SPAN ports emit copies of selected traffic flows, usually for diagnostic purposes.)

Edge Intelligence replaces record servers with record functions on end-user PCs (Figure 5). PCs record, filter and analyze conversations at the edge, close to where the conversations take place. Selected recordings still travel to the central site for archival storage or further analysis, but these recordings are compressed at the edge, yielding a major reduction in network traffic. Business logic, as defined by a manager or supervisor, identifies which conversations to review—the longest, the shortest, calls at certain time of day, calls from certain callers, or other criteria—and downloads those policies to end-user PCs. The PCs filter calls according to the policies and ship back only "calls of interest." The calls travel to the central site as compressed files at low, non-real-time priority, for minimal bandwidth impact. Moreover, PCs can analyze calls themselves and ship back just the results, reducing not only bandwidth utilization but also the demand for central-site analytics servers.

Edge Intelligence gives supervisors more granular control over call selection and filtering than the site-based approach. If desired, supervisors can set different filtering policies for each edge device, not just each record server. Edge Intelligence also leverages PC capacity for a major reduction in server costs, eliminating record servers at distributed sites and avoiding the need for redundant servers since a PC outage affects only one worker. Contention for SPAN ports is eliminated, and less archival storage space is required at the central site, since only selected calls get shipped back.

As with other Edge Intelligence applications, voice recording and analytics scales with the number of call center workers, since each worker's PC provides the necessary processing and storage resources. It also enables valuable new capabilities, like real-time voice analysis and alarms, since there is no network delay between the live conversation and the recording process.

WORKFLOW AUTOMATION

Workflow automation provides another example of the benefits of Edge Intelligence. Workflow, at least for this example, is the series of tasks that make-up any given transaction, including pre-call routing decisions, agent interaction with enterprise applications (or CRM), agent scripting, after-call activities, wrap up, etc. Automating the flow from task to task drives consistency in how transactions are handled, aids in training, eliminates agent keystrokes, and reduces handling time. To provide a highly tailored customer experience based on who is calling and why, a complex contact center may define multiple workflows per site or even per team.

With Edge Intelligence, the policies that apply this business logic reside where the transactions take place—on the agent's PC. Regardless of where those agents reside—at a central call center, a regional site, or a home office—their applications are automatically configured to meet the appropriate requirements—databases accessed, supervisors engaged, automated after-call activities—without affecting other agents and without adding complexity to other applications. For efficiency, common workflows can be created and managed centrally and automatically pushed to PCs at the edge. And because this intelligence resides at the edge, it can be configured uniquely from agent to agent to accommodate special situations. For example, workflow rules on new agents' PCs can take them out of ready status and initiate additional training when call volumes fall below a certain level. Settings on each PC would handle this automatically based on real-time call center data. Moreover, because this intelligence resides at the edge, it can easily be modified to accommodate point-in-time events such as short-term offers or real-time circumstances such as extreme call peaks that necessitate exceptions in call routing.

CONCLUSION

IP communications provides a foundation for applying Edge Intelligence to call centers and other business functions that integrate voice and data. Edge Intelligence distributes application logic to the network edge—closer to the events and information that drive that logic—which can yield significant benefits for today's enterprise call centers:

- Faster, easier application development,
- Increased end-user productivity,
- More reliable system performance,
- Expanded system capacity, flexibility and scalability,
- Reduced capital and operating expense.

EDGE INTELLIGENCE AND SIP

SIP, the Session Initiation Protocol, has become the de facto signaling standard for IP communications, including voice, video, instant messaging and other real-time media. SIP is supported by equipment manufacturers, software houses and service providers worldwide. Like Edge Intelligence, SIP eliminates central control points in favor of intelligent devices at the network edge.

With SIP, an edge device initiates a voice call or other interactive session by sending an INVITE message to another edge device. The called device accepts the INVITE by sending an OK message back to the caller. (Often a SIP proxy server mediates the exchange.) Once the call is accepted, media traffic (e.g. a VoIP packet stream) flows, in general, directly between edge devices, avoiding the bottlenecks caused by intermediate control points.

Thus SIP and Edge Intelligence are complementary. Both facilitate distribution of application processing and control to intelligent edge devices. And both yield more scalable, more robust business services.

ABOUT SPANLINK COMMUNICATIONS

Spanlink Communications is a leading provider of REAL customer interaction solutions that leverage VoIP technology. A Cisco ATP Certified Channel Partner, Spanlink has unmatched experience with Cisco IP Communications solutions and is a leading developer for contact center, collaboration and IP Communications system management solutions for public institutions and medium and large enterprises.

Spanlink customer interaction and system management products are built from the ground up to exploit the benefits of Voice over. Spanlink offers InteractiveEdge products that optimize the way agents and supervisors interact with customers - and with each other - whether they're across the room, in Indiana, or in India. Spanlink® Quality Management is a highly-scalable voice and screen recording and evaluation solution that supports agents and supervisors at any physical location. Cisco Agent Desktop provides agents "anywhere" with the tools they need to manage the contact more efficiently. Cisco Supervisor Desktop promotes team collaboration and coaching for virtual teams.

For more information on Spanlink Communications and our Interactive Edge products for contact centers, visit www.spanlink.com, telephone 800-303-1239 or email Mktg@Spanlink.com.